An Analysis of a Large Scale Habitat Monitoring Application

Robert Szewczyk†, Alan Mainwaring⋆, Joseph Polastre†, John Anderson‡ and David Culler†
†EECS Department, University of California, Berkeley, Berkeley, California 94720
⋆Intel Research Berkeley, 2150 Shattuck Avenue, Berkeley, California 94704
‡College of the Atlantic, 105 Eden St., Bar Harbor, ME 04609

ABSTRACT

Habitat and environmental monitoring is a driving application for wireless sensor networks. We present an analysis of data from a second generation sensor networks deployed during the summer and autumn of 2003. During a 4 month deployment, these networks, consisting of 150 devices, produced unique datasets for both systems and biological analysis. This paper focuses on nodal and network performance, with an emphasis on lifetime, reliability, and the the static and dynamic aspects of single and multi-hop networks. We compare the results collected to expectations set during the design phase: we were able to accurately predict lifetime of the single-hop network, but we underestimated the impact of multi-hop traffic overhearing and the nuances of power source selection. While initial packet loss data was commensurate with lab experiments, over the duration of the deployment, reliability of the back-end infrastructure and the transit network had a dominant impact on overall network performance. Finally, we evaluate the physical design of the sensor node based on deployment experience and a *post mortem* analysis. The results shed light on a number of design issues from network deployment, through selection of power sources to optimizations of routing decisions.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks:**]: Network Architecture and DesignWireless Communications; C.3 [**Special Purpose And Application-Based Systems**]: Real-Time and embedded systems; C.4 [**Performance of Systems**]: Design Studies

General Terms

Performance, Design, Implementation

Keywords

Sensor Networks, Habitat Monitoring, Microclimate Monitoring, Network Architecture, Long-Lived Systems, Application Analysis

1. INTRODUCTION

A broad class of applications are within the reach of contemporary wireless sensor networks (WSNs). These applications share a common structure, where fields of sensors are tasked to take periodic readings, and report results and derived values to a central repository. There are both scientific and commercial applications, for example: microclimate monitoring, plant physiology, animal behavior [16], precision agriculture [9, 3], structural monitoring [4] and condition-based maintenance. These *sense-and-send* applications have widely-varying sampling rates and network bandwidth demands.

In the context of habitat and environmental monitoring, WSNs offer significant advantages. Individual devices can be made sufficiently numerous to take measurements at many locations of interest, and mitigate errors arising from the interpolation and extrapolation from coarser-grained samples. They can be sufficiently small to be co-located with phenomena of interest without altering the parameters to be measured. And they can be unobtrusively embedded in the environment without creating conspicuous landmarks that change the behaviors of its inhabitants.

Long-term unattended operation enables measurement at spatial and temporal scales impractical with human observers or sparsely deployed instruments. The lifetimes made possible with contemporary low-power microelectronics can prolong the duration of experimental observations. At the same time, automation improves the data quality and uniformity of measurement, while reducing data collection costs as compared with traditional human-centric methods. Devices can operate for prolonged periods in habitats that are inhospitable, challenging or ecologically too sensitive for human visitation. Unobtrusive observation is key for studying natural phenomena.

WSNs offer more capabilities than standalone dataloggers and wired instrumentation. Wireless telemetry is valuable because it minimizes observer effects, study site intrusions and environmental alterations. For example, visits to study areas to monitor and download loggers are no longer necessary, while health and status of instrumentation can be monitored remotely. More general networking offers great benefits, such as continuously updated databases of sensor readings accessible through the web, access to live readings from individual sensors, and is key to distributed in-network processing. These capabilities may yield new experimental designs, and paradigms for data publication, dissemination, and scientific collaboration.

We have incrementally deployed several sensor networks of increasing scale and physical extent in a wildlife preserve. While amassing a novel dataset for biological analysis, the annotated data are interesting from a systems perspective. The packet logs from a single-hop and multi-hop network reveal insight on lifetimes, packet yields, network structure and routing. For example, some nodes ran for nearly four months but some for just a few days. Analysis reveals changes in network structure and performance over

SenSys'04, November 3–5, 2004, Baltimore, Maryland, USA.

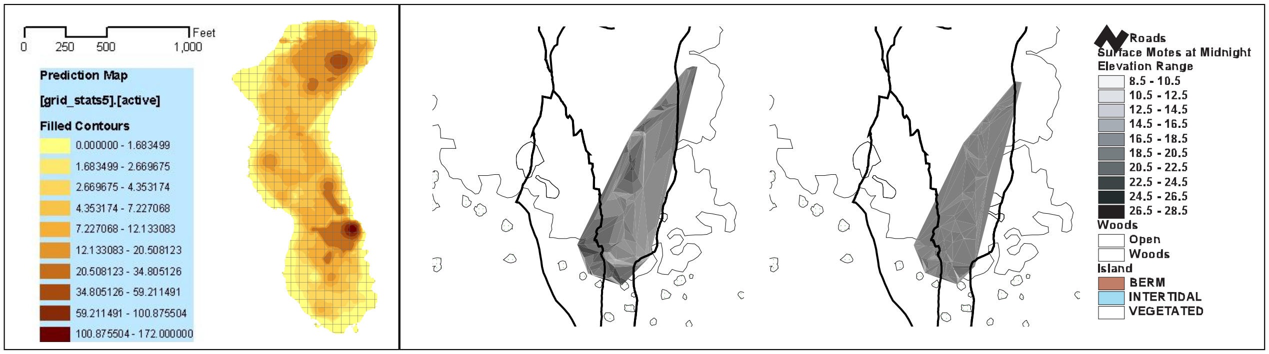

Figure 1: Geospatial distribution of petrels obtained by direct human observation (left) and a particular feature of the habitat (average temperature at midnight in the burrows (center) and on the surface (right) collected from out sensor network)

the lifetime of the deployment. Though the application was simple, it exhibited interesting and unexpected behaviors after its initial setup. Although it is representative of applications with low sampling and bandwidth demands, its architecture and implementation are general and thus provides a reference point for others in this space.

The remainder of this paper presents an analysis of that data collected during the summer and autumn of 2003 from two sensor network deployments on Great Duck Island, Maine. It is organized as follows: Section 2 describes the application, system architecture and realization. Section 3 is an analysis of the data. Section 4 presents experiences and lessons learned. Section 5 discusses related works and Section 6 concludes.

2. SYSTEM

Anticipating to the analyses in Section 3, this section presents background on the application, the system architecture and its implementation. In particular, it describes the tiered network architecture typical of habitat monitoring applications and the design and implementation of its core components.

2.1 Application Background

John Anderson was studying the distribution and abundance of sea birds on an offshore breading colony on Great Duck Island, Maine. He wanted to measure the occupancy of small, underground nesting burrows, and the role of micro-climatic factors in their habitat selection. We hypothesized that a sensor network of motes with appropriate sensors, with TinyOS components for low-power routing and operation, could log readings in a web-accessible database. It seemed plausible that passive infrared (PIR) sensors could directly measure heat from a seabird, or that temperature/humidity sensors could measure variations in ambient conditions resulting from prolonged occupancy. We also wanted, in a modest way, to translate the vision for sensor networks to a concrete reality. This simple application would require creating a complete hardware/software platform, firmly grounded in the needs of a traditional ecological study. It emphasized small mote size, long lifetime, unattended operation, and caused us to consider the verification and ground-truth of sensor readings.

Among the life scientists, Graphical Information Systems (GIS) have become the lingua franca for the visual presentation, analysis and exchange of geospatial data. We imagined a system capable of producing animal density GIS plots like Figure 1, and at the conclusion of 2003, the system could generate such visualizations for micro-climate data. The first GIS plot shows predicted population density on the island based upon direct inspection of burrow occupancy from an entire season of sampling - months of labor resulting in a single plot. The latter two plots show temperatures in the underground burrows and at the corresponding points on the surface. Data for these was collected by our sensor network at midnight on a typical summer evening. Darker colors are warmer temperatures, lighter colors correspond to cooler temperatures, and and shaded area of similar colors are isoplethic temperature regions.

Cooling surface temperatures are apparent, whereas the buffering and insulating properties of the burrows cause them to maintain a nearly constant temperature. Hot-spots in underground burrows are of special interest. For these hot-spots there is mounting evidence from direct inspection and acoustic playbacks that a resident petrel produces the heat. Once the correlation between warmer burrow temperatures and occupancy can be definitively established, expected population density visualizations could be replaced by nightly, or even hourly, sensor data. This would represent a fundamental advancement towards the understanding the distribution and abundance of this species.

2.2 Architecture

The system we deployed has the tiered architecture shown in Figure 2. The lowest end consists of *sensor nodes* that perform communication, computation and sensing. They are typically deployed in *sensor patches*. Depending on the application, they might form a linear transect, a grid region, or a volume of nodes nodes for three-dimensional monitoring. Each sensor patch has a *gateway* that sends data from the patch through a *transit network* to a remote *base station* via the *base station gateway*. We expect mobile *field tools* will allow on-site users to interact with the base station and sensor nodes to aid the deployment, debugging and management of the installation. The base station provides Internet connectivity and database services. It should handle disconnected operation from the Internet. Remote management facilities are a crucial feature of a base station. Typically the sensor data is replicated off-site. These replicas are located wherever it is convenient for the users of the system. In this formulation, a *sensor network* consists of one or more sensor patches spanned by a common transit network and base station.

Sensor nodes are small, battery-powered devices capable of general purpose computation, bi-directional wireless communication, and application-specific sensing. The sizes of nodes are commensurate with the scale of the phenomenon to be measured. Their lifetime varies with duty cycle and sensor power consumption; it can be months or years. They use analog and digital sensors to sample their environment, and perform basic signal processing, e.g., thresholding and filtering. Nodes communicate with other nodes

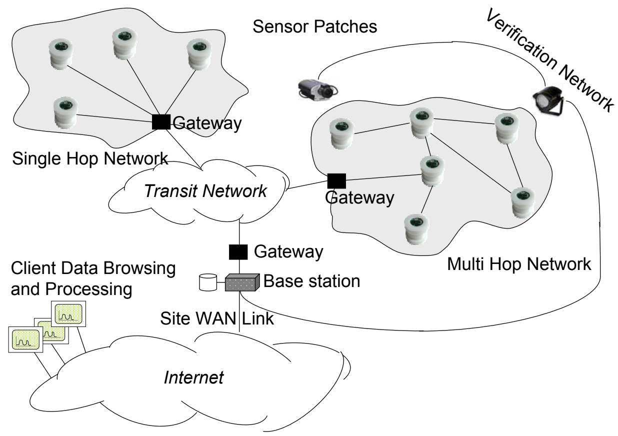

Figure 2: Architecture of the habitat monitoring system

either directly or indirectly by routing through other nodes.

Independent *verification networks* collect baseline data from reference instruments that are used for sensor calibration, data validation, and establishing ground truth. Typically verification networks utilize conventional technologies and are limited in extent due to cost, power consumption and other deployment constraints.

2.3 Implementation

The deployment was a concrete realization of the general architecture from Figure 2. The sensor node platform was a Mica2Dot, a repackaged Mica2 mote produced by Crossbow, with a 1 inch diameter form factor. The mote used an Atmel ATmega128 microcontroller running at 4 MHz, a 433 MHz radio from Chipcon operating at 40Kbps, and 512KB of flash memory. The mote interfaced to sensors digitally using I2C and SPI serial protocols and to analog sensors using the on-board ADC. The small diameter circuit boards allowed a cylindrical assembly where where sensor boards were end caps with mote and battery internal. Sensors could be exposed on the end caps; internal components could be protected by O-rings and conformal coatings. This established a mechanical design where sensor board and battery diameters were in the 1 to 1.5 inch range but the height of the assembly could vary.

We designed two different motes for the application, *burrow motes* for detecting occupancy using non-contact infrared thermopiles and temperature/humidity sensors and *weather motes* for monitoring surface microclimates. The burrow mote had to be extremely small to be deployed unobtrusively in nests typically only a few centimeters wide. Batteries with lithium chemistries were chosen because discharge voltages remained in tolerance for mote, radio and sensors almost to the end of the battery lifetime. We elected not to use a DC boost converter because it introduces noise and increases power consumption. The mote's operation and the quality of its sensor readings depends upon the voltage remaining within tolerance. Should the voltage fall outside the operating range of the mote, radio, or sensors, the results are unpredictable.

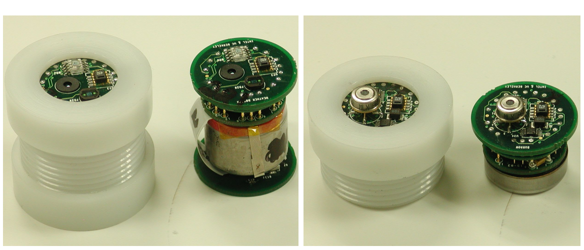

Figure 3: Mote configurations used in the deployment: weather mote (left) and burrow mote (right)

2.3.1 *Burrow and Weather Motes*

Burrow motes monitor temperature, humidity and occupancy of nesting burrows using non-contact passive infrared temperature sensors. They have two sensors: a Melexis MLX90601 non-contact temperature module and a Sensirion SHT11 temperature and humidity sensor. The Melexis measures both ambient temperature ($\pm$ 1°C) and object temperature ($\pm$ 2°C). The Sensirion measures relative humidity ($\pm$ 3.5% but typically much less) and ambient temperature ($\pm$ 0.5°C), which is used internally for temperature compensation. The motes used 3.6V Electrochem SB880 batteries rated at 1Ahr, with a 1mA rated discharge current and a maximum discharge of 10mA. The 25.4mm diameter by 7.54mm tall dimensions of the cell were well suited for the severely size constrained burrow enclosures.

Weather motes monitor temperature, humidity, and barometric pressure. (They also measure ambient and incident light, both broad spectrum as well as photosynthetically active radiation, but these were not used in this application.) They have the following sensors: Sensirion SHT11, Intersema MS5534A barometer, 2 TAOS TSL2550 light sensors, and 2 Hamamatsu S1087 photodiodes. The Intersema measures barometric pressure ($\pm$ 1.5 mbar) and ambient temperature ($\pm$ 0.8°C) used for compensation. The motes used 2.8V SAFT LO34SX batteries rated at 860mAhr, with a 28mA rated discharge current and a maximum discharge exceeding 0.5A. A 25.6mm diameter by 20.3mm height were similar to the Electrochem, permitting a similar packaging technique. The battery exhibits a flat voltage profile for nearly its entire lifetime.

2.3.2 *Mote-based Networks*

In order to conduct viable ecological studies, we need to provide reliable measurements every hour. In both networks we oversampled the environment, and sent the data in a streaming fashion over an unreliable channel. Such approach required maintenance minimal state within the network, while allowing for reconstruction of environmental data at the resolutions required by life scientists.

The first network deployed was an elliptical single hop network. The total length of the ellipse was 57 meters. The network gateway was at the western edge. Nodes in this network performed no routing, they sampled their sensors every 5 minutes and sent results to their gateway. The gateway system was built around two motes communicating over a wired serial link. One of the motes used a TESSCO 0.85dBi omni-directional antenna to interface with the sensor patch. The second mote in the gateway used a Hyperlink 14dBi yagi for a long distance point-to-point link to the base station. At the base station, about 120 meters away, another mote equipped with the yagi antenna received packets from the patch.

The second deployed network was a multi-hop network with a kite-shape to the southwest and a tail to the northeast. Its total length is 221 meters with a maximum width of 71m at the southwest but it narrows to 8m at the northeast. Nodes in this network sampled every 20 minutes and routed packets destined for its gateway. The gateway system configuration was nearly identical to the one used by the single hop network. The sensor patch interface of the gateway periodically sent out routing beacons to seed the network discovery process.

To eliminate the potential interference between the networks, we configured them to operate on different radio frequencies: single-

hop network communicated in 433 MHz band, and the multi-hop in 435 MHz band. Similarly, the long point-to-point links were configured to use the different frequencies (915 and 916 MHz), also non-interfering with patch networks.

2.3.3 *Verification Network*

To understand the correlation between infrared sensor readings from burrow motes and true occupancy, the verification network collected 15 second movies using in-burrow cameras equipment with IR illuminators every 15 minutes. Using a combination of off-the-shelf equipment–cables, 802.11b, power over Ethernet mid-spans, and Axis 2401 camera servers–eight sites were instrumented and five operated successfully. All verification network equipment was physically distinct from the transit and sensor networks with the exception of the laptops at the base station. Scoring the movies by hand in preparation for analysis with sensor data is underway. Evaluation of biological data is not the focus of this paper, so we will not examine the overall impact of the verification network.

2.3.4 *WAN and Base Station*

A DirecWay 2-way satellite system provided WAN connectivity with 5 globally routable IP addresses for the base stations and other equipment at the study site. This provided access to the PostgreSQL relational databases on the laptops, the verification image database, administrative access to network equipment, multiple pan-tilt-zoom webcams and network enabled power strips. A remote server computed the set of database insertions and uploaded the differences every 20 minutes via the satellite link. Were the link unavailable, updates were queued for delivery. The remote databases were queried by replicas for missing data upon link reconnection. Although the upstream bandwidth was small (128Kbps), we did not witness overruns. The base station, satellite link and supporting equipment were powered by a standalone photovoltaic system with an average daily generating capacity of 6.5kWh/day in the summer.

2.4 Media Access and Routing

The network software was designed to be simple and predictable. The radio was duty cycled in our deployments with a technique called low power listening [7]. Low power listening periodically wakes up the node, samples the radio channel for activity, and then returns to sleep if the channel is idle. Packets sent to a low power listening node must be long enough that the packet is detected when the node samples the channel for activity. Once activity is found, the node stays awake and receives the packet, otherwise it goes back to sleep.

The single hop network utilized low power listening but with normal sized packets to its transit gateway. Packets with short preambles can be used because the gateway does not duty cycle the radio – instead it is always capable of receiving packets. The sensor nodes periodically transmitted their sensor readings. They used low power listening to allow the base station to issues commands to change their sample rate, to read calibration data them, and to *ping* the node for health and status information.

The multi-hop network integrated low power listening with adaptive multi-hop routing developed by Woo [20]. Each node selected its parent by monitoring the channel and using the path it expected to be most reliable. Nodes periodically broadcasted their link quality estimates to their neighbors every 20 minutes. This data was used to find reliable bidirectional links. The nodes communicated with each other using low power listening and long packets. We estimated a network neighborhood size of 10 nodes. Given the neighborhood size and sampling rate, we calculated that a 2.2% radio duty cycle would maximize the node's lifetime. We deployed the nodes with these settings and allowed them to self-organize and form the network routing tree.

Table 1: Power profiles for single- and multi-hop deployment. Energy refers to the cost of a single operation. To assess the average power drawn by a subsystem, the cost of a single operation is divided by the period between these operations. Once the rates are set, the average power consumption of the overall application is the sum of power consumed by all subsystems. The battery capacity and average power are neccessary to estimate the projected lifetime.

Subsystem	Energy	Single-hop		Multi-hop	
		period	power	period	power
	(mJ)	(s)	(μW)	(s)	(μW)
Baseline sleep	-	-	56	-	56
Timer	0.0034		62	62	
Incoming packet detection (low power listening)	0.465	1.085	465	0.540	930
Packet transmission (short preamble)	3.92	300	14	-	-
Packet transmission (long preamble)	39.2	-	-	600	64.4
Climate sensing	36.4	300	120	1200	31
Occupancy sensing	35.3	300	118	1200	29
Weather mote (w/o forwarding & overhearing)					
Average power		717		1142	
Expected life (days) (860 mAh battery2.8V)		140		90	
Burrow mote (w/o forwarding & overhearing)					
Average power		714		1141	
Expected lifetime (days) (1000mAh battery3.6V)		127		80	

The software for the burrow and weather motes implement a sense-and-send architecture. Once per sampling period, each mote samples it sensors, combines the readings into a single network packet, and sends it to the base station. Single hop motes sample every five minutes and multi-hop motes every twenty minutes. Each mote listens for incoming packets and dispatches to message handlers upon receipt. When destined for the mote, command interpreter processes the packet, e.g., to change sampling rates or respond to a ping request. Otherwise, the routing subsystem forwards the packet towards its destination.

3. ANALYSIS

This section analyzes the performance of the sensor networks from a systems perspective, considering power consumption, network structure, routing and packet yields. The first deployment started June 8th with the incremental installation of a single hop network. At its peak starting June 16th, the network had 49 motes (28 burrow and 21 weather). A second deployment began July 8th with the incremental installation of a multi-hop network. At its peak starting August 5th, the network had a total of 98 motes (62 burrow and 36 weather). During their combined 115 days of operation, the networks produced in excess of 650,000 observations.

3.1 Lifetime

The design goal was to provide observations for an entire four month field season. We first examine the lifetime of single- and multi-hop motes and compare the achieved performance with estimates. We note that the lifetimes are impacted by the data-collecting

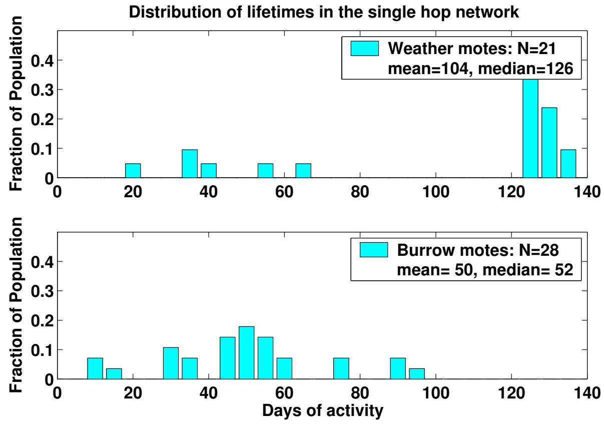

Figure 4: Node lifetimes in the single hop network

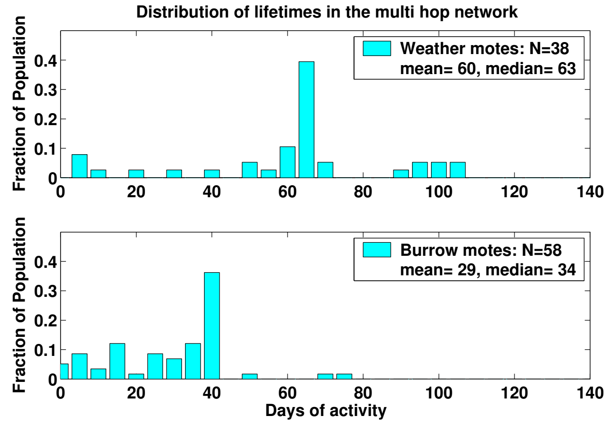

Figure 5: Node lifetimes in the multi-hop network

infrastructure: severe weather (hurricane Isabel) forced the base station to be shut down between September 15 and October 9, and the photovoltaic panels were disconnected for the winter on October 20. In lifetime analysis, we examine the first and last recorded reading, and consequently, we underestimate the lifetime of motes that ceased operation because of base station outages.

The simple structure of the single-hop network makes its power analysis straightforward. Table 1 summarizes the primary power consumption in the system. The original estimates were for 140 days of operation for weather motes and 127 days of operation for burrow motes. In the multi-hop network, motes were expected to last 90 and 80 days, respectively. These estimates were derived without accounting for overhearing or impact of packet forwarding. In addition, we expect the estimates for the burrow motes overestimate the lifetime, since in active state, the power drawn significantly exceeds the battery rating. Such load would decrease the effective battery capacity, though we were unable to quantify that effect prior to the deployment.

Figure 4 shows the observed distribution of mote lifetimes in days in the single hop networks for both weather and burrow motes. We have met the lifetime goals for weather motes – they were operational at the end of deployment, after over 120 days of operation. While some burrow motes operated for over 90 days, the median operation was significantly shorter than our estimate, at 52 days. Two key differences are responsible for this: different power source (heavily taxed by the peak load of the burrow mote) and the harsher operating environment than that of weather motes. We examine the battery performance in more detail below, and analyze the impact of the environment by examining the physical condition of recovered motes in Section 4

Figure 5 shows the distribution of mote lifetimes in days in the multi-hop network. The median lifetime for weather motes is 63 days. It appears at the peak of the histogram, corresponding to motes deployed in July and last reporting on September 15, just before a 23-day gap in data logging. Burrow motes, as in a single hop case, show lifetimes considerably shorter than expected: the median life is 34 days, or under 45% of the original estimate. Nearly 25% of these sensors lasted fewer than 20 days, a mere quarter of the original estimate.

We expect that the shortened lifetime of the multi-hop motes is a direct result of overhearing traffic. In burrow motes, we expect that added load will impact battery capacity; we estimate the overhearing impact from the approximate connectivity and routing graphs. Based on the median lifetime of weather motes, we calculate that the median power draw to be 1600 μW. The estimate discounting the overhearing was 1141 μW, implying that the power draw just attributed to overhearing is over 450 μW. That is 8 times more than the cost of transmission in the multi-hop network! We next examine the battery voltage to determine whether node attrition was a result of depleted batteries.

The batteries used in both burrow and weather motes have flat discharge curves – the battery voltage stays nearly constant over 90% of their capacity, then drops abruptly. Two factors affect the battery voltage sensed by the mote: ambient temperature and mote's current draw. All nodes deployed in this study can sense temperature, so it is possible to normalize the battery voltage readings to a reference temperature. While this temperature-adjusted voltage is a poor measure of remaining capacity, low battery voltage can indicate either excessive current drain or a nearly spent battery.

Figure 6 shows the average voltages reported by motes during their last three hours of operation. A threshold is highlighted on each graph showing the lowest voltage at which a mote will reliably operate. Any device that reported voltages below that cutoff has essentially exhausted its supply, but managed to get a last few reports out. On the other hand, the sharp drop-off in the battery discharge curves implies that a device may not always be able to successfully report once the battery voltage is too low. The voltage thresholds, based on the battery documentation, were selected to be 2.7V for the weather and 3.0V for the burrow motes. We highlight the battery voltage over time for a few selected motes: we plotted the trajectory of the daily voltage for one long lived and one short lived in each mote population. These traces demonstrate examples of both a constant voltage through most of node's life; as well as either a rapid drop in battery voltage or a silent stop at the end of operation.

The clustering of points at particular lifetimes is an artifact of the base station shutdowns discussed above. In the multi-hop networks, nearly 40% of devices are clustered at the lifetimes corresponding to the September shutdown. Since most of the motes form clusters above the threshold voltage (i.e., they still may have remaining energy), we can conclude that the base station outages have had an impact on mean mote lifetime.

Of the original 21 single hop weather motes, 15 were operating at the end of the season on October 20th, and all show relatively high battery voltages. Improper sealing may play a part in shortening the lifespan of the 6 remaining motes.[1] Only 3 of the multi-hop

[1]When we recovered a subset of deployed motes in 2004, 75% of short-lived single-hop weather nodes showed moisture on the

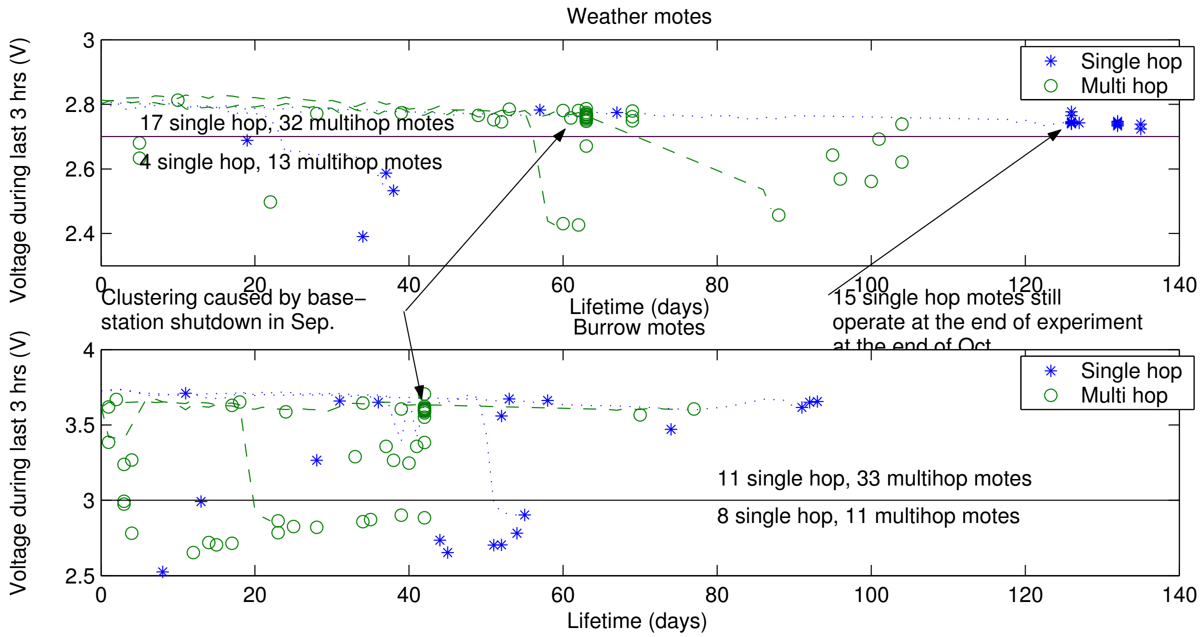

Figure 6: Mote voltage at the end of operation. The cutoff voltages are conservative and have been selected from battery datasheets. The dotted and dashed trails illustrate average daily voltages from representative motes; these are in line with the discharge curves in the datasheets.

weather motes were prematurely terminated because of clear battery problems.

The burrow motes exhibit a similar pattern of end-of-life voltages. In the single hop case, motes drain their battery thoroughly: 8 motes fall below the conservative 3V threshold. We observe a sharp voltage drop at the end of battery capacity – it is possible that the remaining experienced a drop so rapid that it was not recordable. The multi-hop burrow motes exhaust their supply rapidly – we record 7 devices reporting below-threshold voltage readings, and disappearing within the first 20 days. Five other burrow motes stop reporting data within the first 5 days. The battery is heavily taxed to source current for the long preambles in packet transmission and supply energy to overhear the multi-hop traffic in the neighborhood. We attempt to verify the second hypothesis below, in subsection 3.2.

3.2 Multi-hop network structure

Both weather and burrow motes participate in the same multi-hop network, thus we evaluate them together in the context of network topology. Here we evaluate the network structure. The next section examines the dynamic behavior of the routing algorithm.

The low-power listening technique that formed the link layer in the multi-hop network, lowers the cost of listening, while increasing the cost of both transmitting and receiving. Overhearing is costly – there is no early packet rejection, and a packet transmission typically results in the entire one-hop neighborhood receiving the packet. Consequently, the connectivity of the network has an impact on the power consumption of individual nodes. Due to packet size limitations, we did not log the information about neighborhood sizes or packet overhearing; we only recorded partial information about the routing tree (the immediate parent). We approximate the connectivity graph from the parent-child relation: the parent selection algorithm, by design, chooses only symmetric links. For each node, we assume that it can hear *all* parents it chose during the deployment, as well as *all* of the nodes that at any point became its children. To estimate the overhearing cost, we assume that a node hears every transmission within its neighborhood, that leads to some overestimation of heard traffic. On the other hand, the approximation disregards all sibling relations, and accounts only packets that were successfully delivered to the base station. We expect the balance of these feature to result in overhearing estimate that falls below the actual cost.

Mesh networks are attractive because they typically offer many redundant routes. For our purposes, we define a network topology to be robust if a node has more than a single route to the patch gateway. Any node that chooses more than a single parent thought its

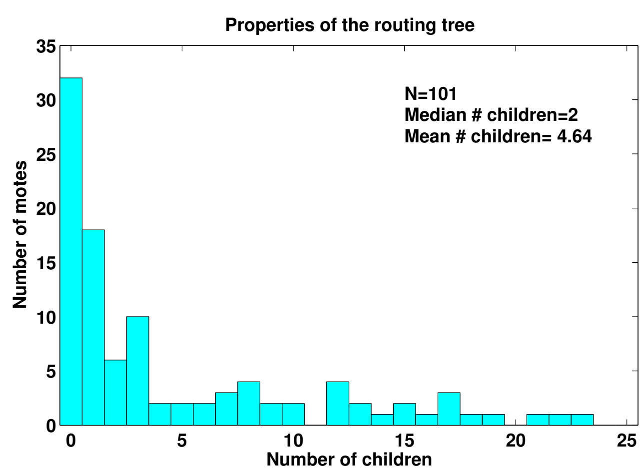

Figure 7: Distribution of children in the routing graph.

---

inside of the package, vs. 25% of long-lived nodes. Unfortunately, for other devices, there was no correspondence between longevity and water exposure.

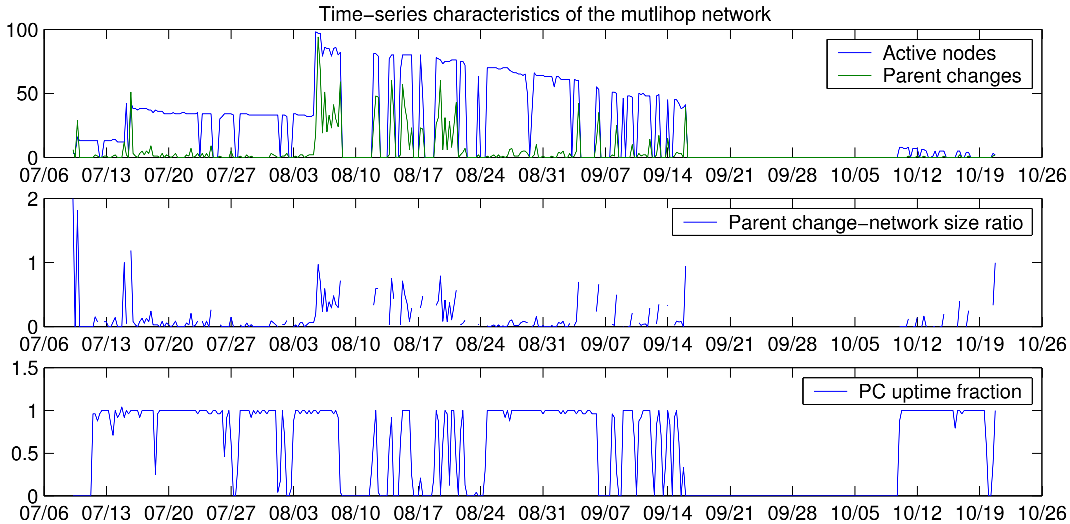

Figure 8: Multi-hop network stability over time. The top graph shows the size of the multi-hop network. Parent change rate is shown in the middle figure. The bottom graph shows the state of the base station logging the data. The base station exhibits a number of outages that impact both the observed network size and the calculation of parent change rates.

lifetime has redundant paths. Only 9 motes out of 101 chose only a single parent; in 6 cases that parent was the gateway to the transit network. The average number of parents for a node is 5.8. We conclude that the deployment placement results in a robust routing tree.

When all communication is symmetric (from the routing perspective), availability of many parent choices implies high rate of overhearing. On the other hand, only the nodes that actually route traffic for others need to listen continually for incoming traffic. As shown in Figure 7, a large portion of the network – 32% – consisted of leaf nodes that never route any packets. These nodes present a clear opportunity for substantial optimization. A leaf node does not need to generate routing beacons – such optimization cuts the originating traffic and decreases the traffic within the one hop neighborhood. A more advanced system, that actually rejects forwarding packets would be even more beneficial.

Burrow motes are in fact much more energy limited than the weather motes, and under most circumstances they should behave as leaf nodes. The deployment did not specifically enforce this, and consequently 48 of the burrow motes were used as parents at some point during their operation; on average these motes routed 75 packets during their lifetimes, with maximum of 570 packets. While the packets routed through the burrow motes were a small fraction of the overall traffic (3600 of 120000 packets), preventing these nodes from routing traffic would have been a simple optimization. Over its lifetime, an average burrow mote overheard nearly 17000 packets (by contrast, the longest lasting burrow mote sourced only about 5300 observations), a significant burden on the limited power supply. Overhearing reduction, while more complex to implement than restricted route selection, needs to be considered in deploying nodes on a limited energy budget.

3.3 Routing Stability

Lab measurements have documented routing stability over periods of hours. We evaluate the stability over weeks and months in a real world deployment. Previously published results show the random fluctuations in link quality world. In addition, the incremental installation of nodes as well as attrition contribute to parent switching.

We begin by looking at the lengths of parent-child relationships within the routing tree. Figure 9 shows a CDF of both the links and the packets delivered over them. Link longevity is measured as a number of consecutive packets successfully delivered through a particular parent. We only count packets sourced rather than routed through the link since packets are sourced at a fixed rate. Because of great range in the link longevity, it is appropriate to plot it on a logarithmic scale. Most links are short lived: the median link is used to deliver only 13 packets. However, most packets are trans-

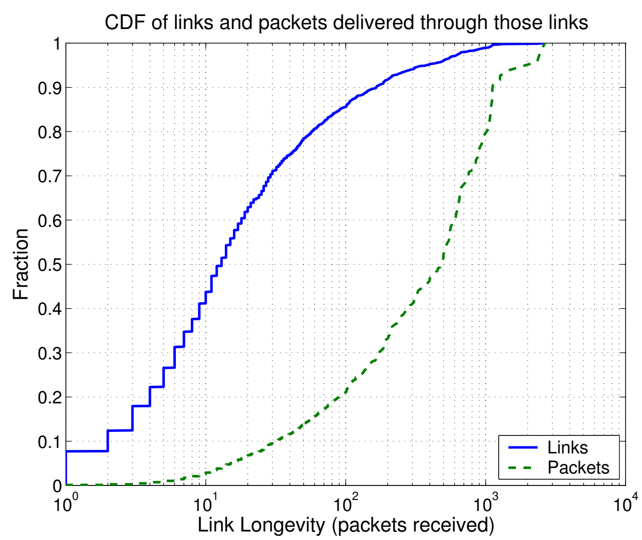

Figure 9: CDFs of parent-child relationship lengths and packets delivered through those links. Long-lived, stable links (ones that delivered more than 100 packets) constitute 15% of all links, yet they are used for more than 80% of the packets.

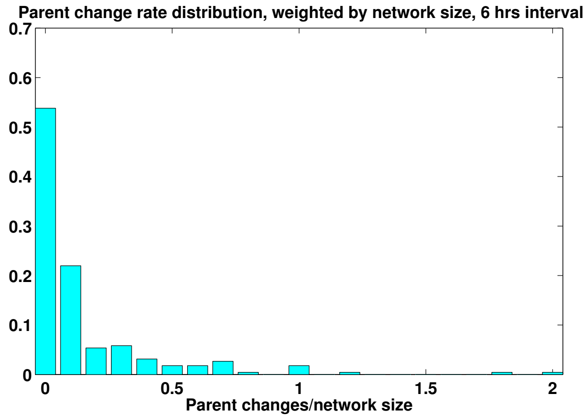

Figure 10: Distribution of parent change rates

mitted over stable stable links. We observe an 80-20 behavior: 80% of the packets are delivered over less than 20% of all links. These stable links last for more than 100 packets – or more than a day and a half. While the distribution may be skewed by nodes that communicate directly with the root, it still can be used to choose the beaconing rates for updating routes.

Another way to look at the stability of the tree is to look at the number of parent changes per time window. Because of the fluctuating network size, we normalize the number of parent changes by the number of motes active over that particular window. Window size affects the analysis and we chose a window of 6 hours, longer than the median longevity of the link. Figure 8 offers a time series view of the multi-hop network. In order to understand its stability over time, we look at two related variables: network size and quality of data logging. Recall that the multi-hop network was installed in 3 stages, concluding with a deployment of burrow motes on August 5. Prior to burrow mote installation, we see a stable mote population. The parent change rate spikes rapidly after the installation of each mote group installation, but settles quickly. After the initial parent turnover, the network is stable; this matches the stability results in [20]. After the deployment of burrow motes the parent change rate remains high at 50% of the node count for a week. This behavior is likely caused by a set of nodes choosing between a few equally good links. The behavior is not directly caused by changes in population size – a period at the end of August corresponds to a similar mix of motes is accompanied my motes disappearing, and yet the parent change rate is nearly 0.

Figure 10 shows the distribution of change rates. Over 55% of time intervals corresponded to times with no changes in the links; 75% experienced less that 0.1 parent changes per 6 hour interval. In a few intervals, the entire network changed routing topology.

3.4 Packet Delivery Effectiveness

Now we examine the effectiveness of packet delivery mechanisms. Recall that both single- and multi-hop networks used the streaming data architecture and relied on oversampling (3x). Previous studies using the same multi-hop routing algorithm reported a 90% packet yield across a 6-hop network [20]. We study the packet delivery over both the short and long term to determine whether the streaming data approach, with no acknowledgments or retransmissions, is sufficient. We note that the networks were operating at a very low duty cycle and that collisions or network congestion should be insignificant.

Figures 11 and 12 show the packet yields from the 4 different kinds of motes. Figure 11 shows the packets delivered to the base station during the first full day of deployment of each network. The results meet expectations set in the indoor lab environment: the weather motes, on average, deliver well over 70% of the packets. The single-hop burrow motes deliver similar yield. The multi-hop burrow motes perform worse (with a median yield of 58%) but within tolerance: the application oversampled the biological signal by 3x. Figure 12 plots the CDF of the packet yields for each mote over every day that mote was active, i.e., delivered a packet. The results are satisfactory for the single hop: the median yield still re-

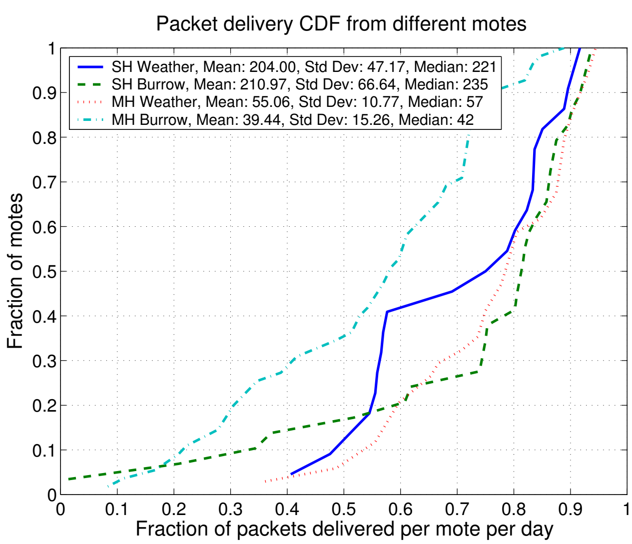

Figure 11: Packet delivery CDF on the first day of complete deployment of the single-hop (June 18, 2003) and the multi-hop network (August 6, 2003). multi-hop weather motes had a median packet delivery of 42 packets (58%). All other motes achieved a median packet yield of over 70%.

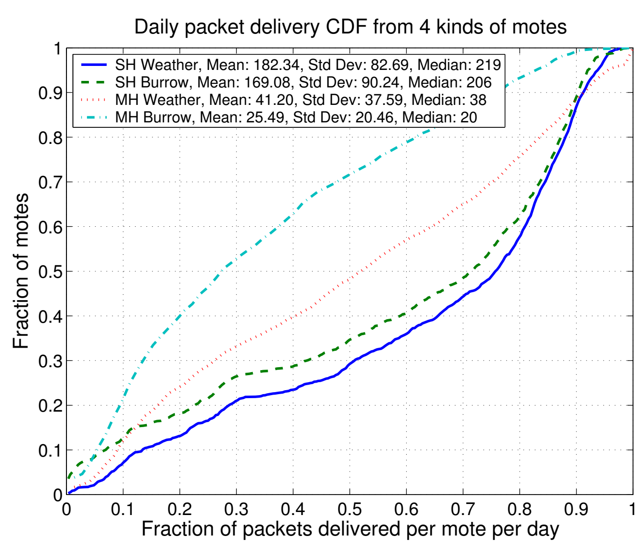

Figure 12: Daily packet delivery CDF over the entire length of the deployment. Motes in the single-hop network deliver a median yield of 70% of packets. The multi-hop network fares much worse, with multi-hop burrow nodes delivering a median yield of just 28%.

mains over 70%. In contrast to Figure 11, the distribution contains many mote-days with substandard performance. On the first day, no mote delivered fewer than 35% of its packets; over the course of the deployment nearly a quarter of the motes performed that badly. The situation is worse in the multi-hop network: the weather motes deliver a median yield of a mere 28%, a performance that jeopardizes the required delivery rates. Some portion of this is due to partial day outages at the base station, but other factors are involved. For example, in the time series in Figure 8, we observe the entire multi-hop network disappear for at least 6 hours (*e.g.*, August 29, and a number of occasions between October 9 and 20). These losses affect the entire network (even the nodes directly communicating with the patch gateway!). The underlying reason for such a correlated outage is likely to correspond to either transit network problems or gateway failures.

To quantify the impact of correlated losses within the multi-hop network, we model the packet losses. The multi-hop network performs no retransmissions. Individual links are lossy and we expect them to deliver similar packet rates to the single-hop network. In addition, the nodes are subject to correlated failures. The combination of these two loss processes results in a delivery rate that is an exponentially decaying function of mote depth.

Figure 13 shows packet yield for each multi-hop mote as a function of its average depth in the routing tree, weighted by number of packets. If we assume that the base station and transit links behave the same as the patch network, the packet yield P could be modeled as l^d where l is link quality and d is a packet depth. The best fit link quality l is 0.72 and the mean squared error is 0.03, for both weather and burrow motes. This result is consistent the mean packet delivery in the single hop network in Figure 11 as well as the link quality data reported in [20] (0.7 and 0.73 for weather and burrow, respectively). It is better than the mean packet yield over the lifetime of the single hop network. A more realistic packet yield model includes the correlated outages that affect the network regardless of the depth. Under these assumption, the model takes a form Al^d, where A corresponds to that deterministic loss for all motes. The best fit parameters curves from this model are shown in Figure 13. The MSE is 0.015 for burrow motes and 0.025 for weather motes. The average link quality estimate of nearly 0.9 shows that the routing layer picks high quality links but the deterministic loss A is also high: 0.57 and 0.46 for weather and burrow motes respectively. These parameters indicate that the best way to improve the packet yields would be to focus on the depth-independent delivery problems.

We conclude that the best-effort delivery in the networking layer was sufficient in the single hop network but not the multi-hop case. Lab experiments typically focus on just the multi-hop network performance; the data from the deployment suggests that base station, transit network and patch gateway have a tremendous impact on a deployed application over its lifetime. Packet delivery needs improvement in the system-wide sense. A communication protocol that addresses independent, local losses (like link-level retransmissions), but ignores the wider spread, correlated outages is not sufficient. End-to-end or custody-transfer approaches may be necessary to address reliable packet delivery.

4. DISCUSSION

This section discusses insights we have gained from the data analysis as well as deploying and working with sensor networks in the field. In some cases, we recommend functionality for future systems. In others, we talk about our struggles with primitive tools for the on-site installation and remote monitoring of sensor networks.

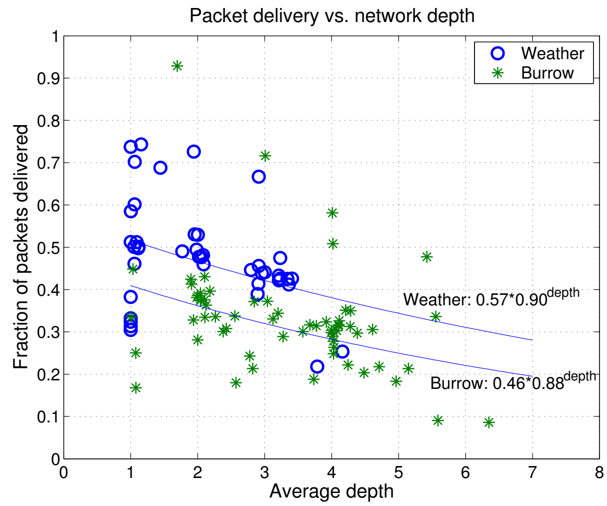

Figure 13: Packets delivered v. average depth. We model packet delivery as a function of the Al^d, where l is link quality, d is the average depth in the routing tree, and A represents packets lost for reasons unrelated to multi-hop routing, like base station or transit network outage

4.1 Node Reclamation

Reclamation is an important practical issue: though researchers often talk about disposable sensor networks, the cost, size, and pollution impact of the devices make reclamation an important final step of an application deployment. Because burrow and weather motes are small, inconspicuous devices, they were easily misplaced in the field. Even with GPS locations, motes deployed in spring were difficult to find when overgrown by summer vegetation. In our deployment, survey flags also identified their positions. Armed with GPS coordinates and flag aids, we launched a recovery effort during the summer of 2004. Of 150 devices deployed in the previous year, we recovered 30 burrow motes and 48 weather motes. The rest of the devices were unrecoverable: either they were moved by animals or their recovery would disturb an active habitat. Examination of the recovered devices allows for a *post mortem* critique of the application and a more direct evaluation of the packaging.

As sensor networks grow in size, we expect the node reclamation to become a crucial part of application planning and deployment cycle. We expect that field tools will aid the recovery process in several ways. By integrating with GPS and accessing the location data, the tools will be able to guide the scientist in the field to a last known location of a mote, even if this mote is no longer functioning. When the network is operating and is actively providing localization service, the field tool should integrate with that service. Finally, the tool could provide direct support for RF direction and range finding (*e.g.*, a directional antenna and an RSSI readout) to guide the scientist to the functioning device in the absence of a more sophisticated localization service.

4.2 Physical Design

There were several interdependent constraints on packaging and power. The mica2dot form factor was predetermined. We wanted a waterproof enclosure design with an internal battery compartment that was small enough for the burrows yet expandable to hold a larger battery in the weather motes. From these constraints, the remaining packaging and power issues and solutions followed. In retrospect, correct approach would be to co-design boards and enclosures with respect to their specific suite of sensors and known environmental conditions.

We produced cylindrical enclosures with threaded end caps from off-the-shelf stock plastic rods on a lathe. In retrospect, more sophisticated enclosures may have been advantageous and more cost effective for several reasons. First, the enclosures required adhesive sealant in addition to O-rings to form a watertight seal; this made both assembly and disassembly time consuming. Second, both the mica2dot and the initial enclosure design assumed an internal antenna. When no internal geometry would provide sufficient communication range, we decided to modify the enclosure to accommodate a standard whip antenna. This design choice complicated assembly (the antenna connector needed to be lined up with the hole absent mechanical guides) and compromised the packaging integrity (we joked that the last step in creating a watertight seal is drilling a hole for the antenna). A new design could properly integrate an antenna into the package or onto the PCB. Third, although the package did not require tools to assemble, the screw-on end caps could not be as tightly fitted as alternative designs that use screws. Fourth, the enclosure dimensions limit the choice of batteries to esoteric cells. A new custom design could make a broader range of less expensive, high capacity cells available. Finally, lack of externally visible signals (like LEDs) made it difficult to immediately verify the liveness of the device – on several occasions we found it necessary to disassemble a device just to turn it on. Overall, we found that the packaging choices hamstrung our ability to deploy individual motes quickly.

The motes recovered in the summer 2004 gave us an opportunity to directly evaluate the packaging effectiveness. External antennas were exposed to wildlife and directly demonstrated risks associated with exposed wires: of the 78 motes, only 13 had fully intact antennas, 6 antennas showed bite marks, but were otherwise unaffected. The remaining antennas were either shortened or removed by animals. The package did not provide complete weather-proofing: 6 burrow motes (20%) and 11 weather motes (22%) had visible droplets of water inside when their enclosures were opened. On the other hand, the breach of package integrity corresponded to shortened lifetime only in single-hop weather motes, which were the first type of devices we deployed. A more sophisticated package might include sensors inside the enclosure to detect condensation; such solution is being employed by researchers at UCLA in the ESS deployment.

For sensor networks to scale to hundreds and thousands of nodes, motes can be touched just once. Assembling a mote, programming it in a test fixture, enclosing it in a package, positioning it the field, and acquiring its survey GPS location is impractical for large numbers of motes. Even the end user who receives a complete, pre-assembled mote ready for deployment faces usability problems of scale. Ideally, the mote should be a completely encapsulated, or packaged, with only the sensors exposed, a non-contact on/off switch, and robust network booting and reprogramming. Issues of programming and network construction should be handled with tools that operate on aggregates of nodes rather than individuals wherever possible.

4.3 Mote Software

The shorter lifespan of burrow motes in the multi-hop network was surprising - nearly 50% less than expected as shown in Figure 4. We were unable to definitively determine the root cause of burrow mote failures although were able to identify a few potential factors. In this section we identify solutions that may assist in future deployments with root cause analysis.

Power monitoring: When using batteries with a constant operating voltage, such as the lithium batteries used in our deployment, battery voltage does not indicate how much capacity is remaining. A more adequate measure of how much work has been performed by the node is needed to calculate each node's expected lifetime. Since our deployment, we have implemented energy counters at the MAC layer and in our sensing application. Each counter keeps track of the number of times each operation has occurred (*e.g.*, sensing, receiving or transmitting bytes, total amount of time the CPU is active). By keeping track of this data, nodes can report on the number of packets forwarded or overhead. A more complete and automated approach to power profiling is described in [15] We can improve our lifetime estimate through additional health information from each mote. Metrics, such as these energy counts, are crucial to predicting the performance of the deployed network.

Integrated Data-logging: In the current system, packets are sent without any acknowledgments or retransmissions. Sensor readings can be lost at any level of the network from the node to the base station gateway. However, nodes at each level of the network could log data into local non-volatile memory as it heads towards the base station.

A protocol can attempt to incorporate data logging with reliable message delivery or custody transfer protocols to reduce data loss. For example, given 512KB of local storage, 64 byte sensor readings, and a sampling interval of 20 minutes, each node could store 113 days of its own readings. This buffering could mitigate downtime at the added expense of buffering packets along the path. The buffering is free unless the packet is written to flash memory; writing flash memory is approximately four times more costly than sending the message.

Integrated logging allows nodes to retain data when disconnections occur. A node may be disconnected from another mote (such as a parent), the gateway, the base station gateway, one of the base stations, or the data base service. Since large scale deployments are relatively costly, it may be worth taking measures to retain the data if the reduction in longevity can be tolerated.

Finally integrated logging allows for a much more thorough *post mortem* analysis. The information contained in those logs could allow a sensor network developer to analyze the network neighborhoods and topology over time, as such it would be an invaluable debugging and analysis tool.

4.4 External Tools

Currently a great deal of expertise is required to install these networks; rather we would prefer to enable the specialist and the non-specialist alike to accomplish that easily. We identify two levels of tools that provide assistance when deploying sensor networks: field tools that run on small, PDA-class devices and client tools that run on larger laptops class machines situated either at the local field station or more distantly at businesses or universities many miles away. For each class, we note the functionality that would be useful from our own experiences in the field. Additionally we would like to stress the utility of backend analysis and visualization tools being available from the start of the deployment.

Field Tools Functionality:

1. Run self-check. Before placing a mote in the field, where is may not be touched again for months or years, it should be possible to run a final self diagnostic to verify the mote's health. If the device is healthy, it can be left alone, but otherwise repairs must be taken.
2. Show network neighborhood. While incrementally deploying a network in the field, oftentimes one needs to see whether a mote has been placed within range of the rest of the network. Placement is often guided by non-networking factors, e.g., factors of biological interest.
3. Show network statistics. While within range of the mote, it

can be useful to query it for basic packet level statistics. Once a mote has joined a network, it can be useful to monitor how well its networking subsystem is operating before moving onto new locations.

Client Tools Functionality:

1. Re-task nodes (reprogram the sensor network). From time to time, application software upgrades and patches will become available and it will become necessary to upgrade application software. It should be possible to do this from a field station or remotely on the Internet.
2. Show who-can-hear-whom relationships. Show the graph estimating the radio neighborhoods around each node as well as the routing tree currently in use for the network. This is typical of the diagnostic information that technicians would need to monitor network operation.
3. Show when mote lasted reported in. Report when each mote was last heard from. This is a very common and very useful statistic for the layperson looking for signs of a malfunctioning mote or network.

The usefulness of these tools is suggested by a large number of scenarios that arose in the study area. Although the spatial distribution of nodes was driven by the interests of our biologist, these tools can show the density of each level of the multi-hop network. A simple GUI could display when each node was last heard from. An optional alarming mechanism to notify on-site staff when a node failed to report is needed functionality for non-technical on-site staff.

Backend Analysis and Visualization: The back-end infrastructure such as the transit network, base stations and relational databases were deployed before the motes so that packet logs were ready to be captured as soon as sensors nodes began to report in. When deploying new motes, it was possible to see their records being inserted into the database and thus know they were alive. This was a primitive but sufficient means of creating a network when combined with the ability to issue SQL queries against the growing database. The queries allowed one to retrieve statistics such as the network size and when motes were last heard from. A web-based time series visualization was used immediately to track the occupied burrows (by ecologists) and to verify sensing accuracy (by computer scientists).

A variety of tools have since been developed for data analysis as well as network monitoring. Ideally, these GUIs, visualizations, and statistical tools should be available at deployment time as well to enrich the suite to client tools that are available. The statistics one performs on a corpus of data, such as lifetime analysis or voltage profiling, may have great practical utility during phases of network construction and early days or weeks of operation as well. Many of the graphs in Section 3 would be of interest to practitioners deploying their own large scale sensor networks out in the field. The MoteView tool from Crossbow is an example of a deployment and monitoring program of this sort.

5. RELATED WORK

Historically, there have been differences between data loggers and wireless devices like burrow and weather motes. Most importantly, loggers lacked networking and telemetry. Although some vendors have development environments for OEMs, data loggers have been primarily turn-key, rather than flexible devices with open programming environments. Data loggers, such as the Hobo [13], can be larger and, some would argue, more expensive. Some models require external wiring to adjacent equipment. Previously, due to size, external wiring, and organism disturbance, data loggers were found to be unsatisfactory for use on the island.

Vendors such as Campbell Scientific and Onset Computer Corporation now offer radio telemetry systems [1, 12]. This provides wireless connectivity to remote loggers for daily data offload. These systems support potentially hundreds of loggers, with repeaters extending the multi-kilometer link distance. The ad hoc networking within sensor networks remains a distinguishing feature.

Other habitat monitoring studies used one or a few sophisticated weather stations an “insignificant distance” from the study area. With this method, biologists cannot gauge whether a weather station, for example, actually monitors a different micro-climate due to its distance from the organism being studied. A few widely-spaced instruments may give biologists a distorted view of local phenomena. Instead, we wanted to enable monitoring on the scale of the organism, in our case a bird, and the microclimates of distinct nesting areas [5, 17].

Habitat monitoring with WSNs has been studied by others. Cerpa et. al. [2] propose a multi-tiered architecture for habitat monitoring. The architecture focuses primarily on wildlife tracking instead of habitat monitoring. A PC104 hardware platform was used for the implementation with future work involving porting the software to motes. Work with a hybrid PC104 and mote network has been done to analyze acoustic signals [19]; long term results and reliability data may be pending. Wang et. al. [18] implement a method to acoustically identify animals using a hybrid iPaq and mote network.

GlacsWeb [11] is a system that exhibits characteristics at the intersection of the WSN and telemetry augmented datalogging. The overall architecture is similar to GDI system architecture. The devices within a patch (GlacsWeb Probes) are a design point close to motes; the transit network is very reminiscent of the commercial datalogging equipment, both in hardware capability (*e.g.*, powerful radios) and in software design (store and forward architecture, daily scheduled communication).

ZebraNet [10] is a WSN for monitoring and tracking wildlife. ZebraNet nodes are significantly larger and heavier than motes. The architecture is designed for an always mobile, dynamic, multi-hop wireless network. In most respects, this design point is significantly different from our domain of stationary sensor network monitoring.

At UC James Reserve in the San Jacinto Mountains, the Extensible Sensing System (ESS) monitors ambient micro-climate below and above ground, avian nest box interior micro-climate, and animal presence in 100+ locations within a 25 hectare study area. Individual nodes with up to 8 sensors are deployed along a transect, and in dense patches, crossing all the major ecosystems and environments on the Reserve. The sensor data includes temperature, humidity, photosynthetically active radiation (PAR), and infrared (IR) thermopile for detecting animal proximity.

ESS is built on TinyDiffusion [6, 14] routing substrate, running across the hierarchy of nodes. Micro nodes collect low bandwidth data, and perform simple processing. Macro sensors organize the patches, initiate tasking and process the sensor patch data further. They often perform functions of both cluster heads and patch gateways. In case of a macro sensor failure, the routing layer automatically associates macro sensors with the nearest available clusterhead. The entire system is time-synchronized, and uses SMAC for low power operation. Data and timestamps are normalized and forwarded to an Internet publish-and-subscribe middleware subsystem called Subject Server Bus (SSB), whereby data are multicast to a heterogeneous set of clients (e.g., Oracle, MatLab, and LabVIEW) for processing and analysis of both historical and live data streams.

ESS makes an aggressive use of hierarchy within a patch; the diversity of sensors can also be used for verification of data. The SSB is a noteworthy departure from the architecture in Figure 2 – it allows for natural integration of triggered features into the system in addition to data analysis.

California redwoods are such large organisms that their life cycle can be measured through microclimate observations. Having developed models for their metabolism, biologists are now using sensor networks to verify and refine these models. The sensor network measures direct and incident photosynthetically active radiation (PAR), temperature, and relative humidity. In the fall of 2003, 70 nodes were deployed on a representative tree in the middle of the forest, reporting data every minute. Biologists intend to grow the network to both interior and edge trees in a grove.

The network collecting this information is an instantiation of the Tiny Application Sensor Kit (TASK) [8]. The macro sensors in the patch run a version of TinyDB query processing engine that propagates queries and collects results from a multi-hop network. There is no separate transit network – the patch bridges directly to the base station. The base station runs a TASK server that logs data, queries and network health statistics. TASK server is capable of running on a macro sensor. Deployment and in the field debugging are aided by a PDA-class device running a field tool, that allows for connectivity assessment and direct querying of individual sensors. To achieve low power operation, the entire network is time-synchronized and duty-cycled. TASK has a health query as one of the options, which could obtain voltage, routing, neighborhood and other networking state that could facilitate the analyses in Section 3.

6. CONCLUSIONS

We have presented the system architecture, implementation, and deployment of two wireless sensor networks. These networks used a total of 150 nodes in both single-hop and multi-hop configurations. During four months of operation, they compiled a rich dataset with more than 650,000 records that are valuable both to biologists and computer scientists. This paper analyzed this data from a systems perspective of lifetime, network structure, routing and yield.

Lifetime estimates based on micro-benchmarks accurately predicted the single-hop mote longevity but diverged for multi-hop motes. We believe that moisture penetrating enclosures caused attrition in both networks. This obscures the fundamental roles that overhearing and routing packets have on longevity. A sense-and-send application design and best-effort routing layer was sufficient for the single-hop network. About 50% of the losses in the multi-hop routing tree, can be modeled as a deterministic factor; the remaining losses show the expected exponential decay of yield as a function of depth. Each of the packet loss components represents area for improvement, and may need to be addressed differently – custody transfer mechanisms may be a suitable way to handle the deterministic losses, while link-level retransmissions would eliminate the depth-dependent loss.

Motes with smaller battery capacities were observed routing on behalf of motes with larger battery capacities. While routing could be done preferentially by motes with larger batteries, reducing the burden of large packet preambles represents a larger potential savings. Time synchronized low power listening stacks, such as the ones being developed by Crossbow, begin to address this by synchronizing senders and receivers sufficiently to allow receivers to remain asleep during more of the preamble reception time, awaking slightly before the arrival of the packet payload.

After the initial setup, the network exhibited stable periods of organization punctuated by periods of increased re-parenting rates. This was caused primarily from the incremental addition of motes. In general, more than 80% of the packets were routed using less than 20% of the links. Across quantized 6-hour time intervals, in over 50% there were no re-parenting in the network, and in more than 75% there were less than 10% re-parentings. Base station availability, which reflects the reliability of the Windows laptops, background services for logging data into Postgres, and photovoltaic power was lower than expected and resulted in lost data.

We have another year of experience with the practical challenges of deploying sensor networks in the field. We have identified the need to eliminate the logistical overheads required to install each mote, and the need for field tools for in-situ installation and monitoring. We discussed the potential value of logging sensor readings to local storage to allow post-deployment reclamation of data that was not successfully sent to the base station. While not a replacement for reliable networking, it may be appropriate for deployments where motes can be reclaimed in good working condition. Reclamation in our case was complicated because motes were often moved and buried by animals.

Working indoors and in the field allows sensor network researchers to compare laboratory results with observations from the real world over a period of months. In a number of cases (*e.g.*, lifetime prediction for the single-hop network), the micro-benchmarks transfered well into a deployed application. In other cases (*e.g.*, the lower packet yields from the multi-hop network) we saw an interesting departure from the expected behavior. Often these new behaviors emerged as the network was reconfigured, expanded, or re-deployed; others were an artifact of aging. Many of those behaviors cannot be observed without building complete systems and deploying them in realistic conditions.

Acknowledgments

The authors wish to thank the many contributors and supporters of this work: John Anderson, Andrew Peterson, Steven Katona, Jim Beck, Phil Buonadonna, Eric Paulos, Wei Hong, David Gay, Anind Dey, Earl Hines, Brenda Mainland, Deborah Estrin, Michael Hamilton, and Todd Dawson. This work was supported by the Intel Research Laboratory at Berkeley, DARPA grant F33615-01-C1895 (Network Embedded Systems Technology), the National Science Foundation, and the Center for Information Technology Research in the Interest of Society (CITRIS).

7. REFERENCES

[1] CAMPBELL SCIENTIFIC. Spread spectrum radio modems. `ftp://ftp.campbellsci.com/pub/outgoing/lit/b_rf400.pdf`, July 2004.

[2] CERPA, A., ELSON, J., ESTRIN, D., GIROD, L., HAMILTON, M., AND ZHAO, J. Habitat monitoring: Application driver for wireless communications technology. In *2001 ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean* (San Jose, Costa Rica, Apr. 2001).

[3] FAUCHALD, N. Wi-fi hits the vineyard. `http://www.winespectator.com/Wine/Daily/News_Print/0,2463,2539,00.html`, July 2004.

[4] GLASER, S. D. Some real-world applications of wireless sensor nodes. In *Proceedings of SPIE Symposium on Smart Structures and Materials/ NDE* (San Diego, CA, USA, Mar. 2004).

[5] HAPPOLD, D. C. The subalpine climate at smiggin holes, Kosciusko National Park, Australia, and its influence on the

biology of small mammals. *Arctic & Alpine Research 30* (1998), 241–251.

[6] HEIDEMANN, J., SILVA, F., AND ESTRIN, D. Matching data dissemination algorithms to application requirements. In *ACM SenSys 2003* (Los Angeles, CA, USA, Nov. 2003), ACM Press.

[7] HILL, J., AND CULLER, D. Mica: a wireless platform for deeply embedded networks. *IEEE Micro 22*, 6 (November/December 2002), 12–24.

[8] HONG, W., MADDEN, S., FRANKLIN, M., AND HELLERSTEIN, J. TASK: Tiny Application Sensor Kit. `http://berkeley.intel-research.net/task/`, 2004.

[9] INTEL CORP. `http://www.intel.com/labs/features/rs01031.htm`, 2004.

[10] JUANG, P., OKI, H., WANG, Y., MARTONOSI, M., PEH, L.-S., AND RUBENSTEIN, D. Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with ZebraNet. In *Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)* (San Jose, CA, USA, Oct. 2002), ACM Press, pp. 96–107.

[11] MARTINEZ, K., HART, J. K., AND ONG, R. Environmental sensor networks. *IEEE Computer 37*, 8 (Aug. 2004), 50–56.

[12] ONSET COMPUTER CORP. Radio modem and remote site manager software: Wireless communication for HOBO® weather stations. `http://www.onsetcomp.com/Products/Product_Pages/Wireless/wireless_modem.html`, Feb. 2004.

[13] ONSET COMPUTER CORPORATION. HOBO weather station. `http://www.onsetcomp.com`.

[14] OSTERWEIL, E., AND ESTRIN, D. TinyDiffusion in the Extensible Sensing System at the James Reserve. `http://www.cens.ucla.edu/~eoster/tinydiff/`, May 2003.

[15] SHNAYDER, V., HEMPSTEAD, M., RONG CHEN, B., WERNER-ALLEN, G., AND WELSH, M. Simulating the power consumption of large-scale sensor network applications. In *ACM SenSys 2004* (Baltimore, MD, USA, Nov. 2004), ACM Press.

[16] SZEWCZYK, R., OSTERWEIL, E., POLASTRE, J., HAMILTON, M., MAINWARING, A., AND ESTRIN, D. Habitat monitoring with sensor networks. *Communications of ACM 47*, 6 (2004), 34–40.

[17] TOAPANTA, M., FUNDERBURK, J., AND CHELLEMI, D. Development of Frankliniella species (Thysanoptera: Thripidae) in relation to microclimatic temperatures in vetch. *Journal of Entomological Science 36* (2001), 426–437.

[18] WANG, H., ELSON, J., GIROD, L., ESTRIN, D., AND YAO, K. Target classification and localization in habitat monitoring. In *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2003)* (Hong Kong, China, Apr. 2003).

[19] WANG, H., ESTRIN, D., AND GIROD, L. Preprocessing in a tiered sensor network for habitat monitoring. *EURASIP JASP Special Issue on Sensor Networks 2003*, 4 (Mar. 2003), 392–401.

[20] WOO, A., TONG, T., AND CULLER, D. Taming the underlying challenges of reliable multihop routing in sensor networks. In *ACM SenSys 2003* (Los Angeles, CA, USA, Nov. 2003), ACM Press, pp. 14–27.